arXiv:1602.01206v1 [stat.AP] 3 Feb 2016

denoiseR: A Package for Low Rank Matrix Estimation

Julie Josse
Agrocampus - INRIA

Sylvain Sardy
University of Geneva

Stefan Wager
Stanford University

Abstract

We present the R package **denoiseR** dedicated to low-rank matrix estimation. First, we briefly review existing methods including singular value shrinkage and a parametric bootstrap approach. Then, we discuss how to extend the methods to missing values and suggest a general iterative imputation algorithm. It includes an extension of the Stein Unbiased Risk Estimate to missing values for selecting tuning parameters. Finally, we compare and apply the methods with many experiments.

Keywords: Low-rank matrix estimation; singular values shrinkage; bootstrap; SURE; count data; correspondence analysis; clustering; missing values; matrix completion.

1. Introduction

Consider the model where the data matrix X with n rows and p columns is generated as

$$X = \mu + \varepsilon \quad \text{with} \quad \varepsilon_{ij} \overset{\text{iid}}{\sim} \mathcal{N}\left(0, \sigma^2\right), \tag{1}$$

and μ is a matrix of unknown rank k. The statistical aim is to recover the signal from the noisy data. The popularity of model (1) can be explained because both the modelisation "signal+noise" and the low rank assumption seem very plausible to describe a large variety of phenomena. Recently, it emerged in the machine learning community as a powerful way to address the system of recommandation problems such as the famous Netflix challenge (Netflix 2009). A classical solution to recover the signal μ consists in solving

$$\text{argmin}_{\mu} \left\{ \|X - \mu\|_2^2 : \text{rank}\,(\mu) \leq k \right\}, \tag{2}$$

for a given k, which solution (Eckart and Young 1936) is the truncated singular value decomposition (SVD) of the matrix $X = UDV^\top$ at the order k, namely

$$\hat{\mu}_k = \sum_{l=1}^{k} d_l \, u_l \, v_l^\top, \tag{3}$$

where d_l are the singular values organized in decreasing order. The so-called *regularization parameter* k plays a key role and a challenge is to select it from the data. Despite the wide use of estimator (3), other regularization techniques have recently been proposed for the estimation of μ.

These estimators, implemented in the R package **denoiseR**, are described in Section 2. We also detail a recent proposition of Josse and Wager (2014) to go outside the Gaussian framework and denoise count data for instance. Then, in Section 3, we extend the methods to the incomplete case. We tackle the challenging task of selecting the regularization parameter(s) when missing values are present in X, and as an aside, we propose a new method to impute data. Finally, we carry-out many experiments and illustrate the methods on data sets from different fields.

denoiseR is the first package implementing many low-rank matrix estimation methods in the R language. When an implementation is available in another language it will be mentionned.

2. Singular values shrinking and thresholding

2.1. Finite sample

Keeping the first k singular values and setting the others to zero in (3) amounts to hard thresholding the singular values. The most popular alternative consists in soft thresholding them (see Donoho and Johnstone (1994) for definitions of hard and soft thresholding functions):

$$d_l \max\left(1 - \frac{\lambda}{d_l}, 0\right). \tag{4}$$

Singular values smaller than a quantity λ are set to 0 and the others are shrunk towards 0 by an amount λ (note that with hard thresholding, the later are not shrunk towards 0). The estimator associated with (4) is the closed form solution to

$$\text{argmin}_\mu \left\{ \|X - \mu\|_2^2 + \lambda \|\mu\|_* \right\},$$

where $\|\mu\|_*$ is the nuclear norm of the matrix μ. Many studies (Shabalin and Nobel 2013; Josse and Sardy 2015; Josse and Wager 2014) showed with simulations that the soft-thresholding estimator gives small mean squared error (MSE) with low signal to noise ratio (SNR) but struggles in other regimes.

For more flexibility, the adaptive trace norm estimator (ATN) of Josse and Sardy (2015) uses two regularization parameters (λ, γ) to threshold and shrink the singular values:

$$d_l \max\left(1 - \frac{\lambda^\gamma}{d_l^\gamma}, 0\right). \tag{5}$$

This estimator denoted $\hat{\mu}_{(\lambda,\gamma)}$ is the closed form solution to

$$\text{argmin}_\mu \left\{ \|X - \mu\|_2^2 + \alpha \|\mu\|_{*,w} \right\},$$

where $\|\mu\|_{*,w} = \sum_{l=1}^{\min(n,p)} \omega_l d_l$ is a weighted nuclear norm with $\omega_l = 1/d_l^{\gamma-1}$ and $\alpha = \lambda^\gamma$. ATN parametrizes a rich family of estimators and includes (3) and (4) as special cases. As

an insight of its good behavior, we observe that, for $\gamma > 1$, the smallest singular values in (5) are more shrunk in comparison with the largest ones, which is satisfactory since the smallest singular values are responsible for instability.

The parameters (λ, γ) can be selected with cross-validation. However, to avoid such a computational intensive method, Josse and Sardy (2015) extended the results of Candes, Sing-Long, and Trzasko (2013) and define a Stein unbiased estimate of the risk (Stein 1981) $\text{MSE}(\lambda, \gamma) = \mathbb{E}\|\mu - \hat{\mu}_{(\lambda,\gamma)}\|^2$ as follows:

$$\text{SURE}(\lambda, \gamma) = -np\sigma^2 + \sum_{l=1}^{\min(n,p)} d_l^2 \min\left(\frac{\lambda^{2\gamma}}{d_l^{2\gamma}}, 1\right) + 2\sigma^2 \text{div}(\hat{\mu}_{\lambda,\gamma}), \tag{6}$$

where the second term corresponds to the residuals sum of squares (RSS) whereas the last one is the divergence. The expectation of the divergence corresponds to the degrees of freedom. Note that SURE requires to know the noise variance σ^2. Inspired by generalized cross-validation (Craven and Wahba 1979), Josse and Sardy (2015) derived generalized SURE

$$\text{GSURE}(\lambda, \gamma) = \frac{\text{RSS}}{(1 - \text{div}(\hat{\mu}_{\tau,\gamma})/(np))^2}, \tag{7}$$

which does not require knowledge of σ^2. Aiming at good rank recovery, they also employed the quantile universal threshold (QUT) to select λ (Giacobino, Sardy, Diaz Rodriguez, and Hengartner 2015). The rationale of QUT is to select the threshold λ^{QUT} at the bulk edge of what a threshold should be to reconstruct the correct model with high probability under the null hypothesis that$\mu = O$ (the $n \times p$ matrix with zeros for entries). For that specific value λ^{QUT}, they minimize $\text{SURE}(\lambda^{\text{QUT}}, \gamma)$ over γ.

2.2. Asymptotic framework

Other estimators applying non linear transformations of the singular values have been suggested as alternative to (3) in asymptotic frameworks. More precisely, when both the number of rows ($n = n_p$) and columns (p) tend to infinity while the rank of the matrix stays fixed, Shabalin and Nobel (2013) and Gavish and Donoho (2014) showed that the estimator the closest to the true signal μ in term of MSE has the following form, with $n_p/p \to \beta$, $0 < \beta \leq 1$:

$$\frac{1}{d_l}\sqrt{\left(d_l^2 - (\beta - 1)n\sigma^2\right)^2 - 4\beta n\sigma^4} \cdot 1\left(l \geq (1 + \sqrt{\beta})n\sigma^2\right), \tag{8}$$

where 1() is the indicator function. If the noise variance is unknown, the authors suggested

$$\hat{\sigma} = \frac{d_{med}}{\sqrt{n\mu_\beta}}, \tag{9}$$

where d_{med} is the median of the singular values of X and μ_β is the median of the Marcenko-Pastur distribution. Note that Gavish and Donoho (2014) also provided optimal shrinkers using other losses than the Frobenius one, namely the Operator and Nuclear losses.

In another asymptotic framework, considering n and p as fixed, but letting the noise variance σ^2 tend to zero, Verbanck, Husson, and Josse (2015), showed that the best estimator to recover the signal in term of MSE has the following form:

$$d_l\left(\frac{d_l^2 - \sigma^2}{d_l^2}\right) \cdot 1(l \leq k). \tag{10}$$

Each singular value is multiplied by a quantity which can be seen as the ratio of the signal variance over the total variance. This estimator corresponds to a truncated version of the one suggested in Efron and Morris (1972). If the noise variance is unknown, the authors suggested

$$\hat{\sigma}^2 = \frac{\left\| X - \sum_{l=1}^{k} u_l d_l v_l \right\|_2^2}{np - nk - kp + k^2}. \tag{11}$$

Note that the variance estimate requires a value for the rank k, for instance estimated by cross-validation (Josse and Husson 2011).

Josse and Sardy (2015) and Josse and Wager (2014) highlighted that estimators (10) and (8) provide good recovery in their asymptotic regimes, but recommend ATN in other situations.

2.3. Low rank matrix estimation with Gaussian noise

In this section, we use **denoiseR** to estimate a low-rank signal with the estimators listed in the previous sections. After installing the **denoiseR** package, we load it. First, we generate a data set of size $n \times p$ of rank k according to model (1) with the function `LRsim`. The amount of noise is defined by the argument `SNR` and is calculated as $1/(\sigma\sqrt{np})$. This function returns the simulated data in `X`, the signal in `mu` and the standard deviation of the noise in `sigma`.

```
R> library("denoiseR")
R> Xsim <- LRsim(n = 200, p = 500, k = 10, SNR = 4)
```

Then, we denoise the data with ATN (5) using the function `adashrink` as follows:

```
R> ada.gsure <- adashrink(Xsim$X)
R> muhat <- ada.gsure$mu.hat
R> ada.gsure$nb.eigen
[1] 10
```

By default, the tuning parameters (λ, γ) are selected with GSURE (7). The function outputs the estimator of the signal in `mu.hat` which is a matrix of size $n \times p$, the number of non-zero singular values in `nb.eigen` (here 10), the optimal γ in `gamma`, the optimal λ in `lambda` as well as the SVD of the estimator $\hat{\mu}_{(\lambda,\gamma)}$ in `low.rank` and the singular values in `singval`.

It is possible to select the tuning parameters with the other strategies by specifying the argument `method` either at `SURE` or `QUT`. The first one is recommended to estimate the signal μ, and the second to estimate the rank of μ. As detailed in Section 2.1, both strategies require the variance of the noise (as opposed to GSURE).

```
R> ada.sure <- adashrink(Xsim$X, method = "SURE")
Warning message:
In adashrink(Xsim$X, method = "SURE") : sigma estimated by MAD: 0.000811
```

If σ^2 is unknown, `adashrink` calls by default the function `estim_sigma` with the argument `method = "MAD"` to estimate it with (9). If the low-noise (LN) framework seems more plausible, estimator (11) can be obtained by specifying the argument `method = "LN"`, which requires knowing the rank k. Otherwise, k is estimated by default by cross-validation as implemented in the `estim_ncp` function of the **FactoMineR** package (Husson, Josse, Le, and Mazet 2015):

```
R> sigmahat <- estim_sigma(Xsim$X, k = 10, method = "LN")
R> sigmahat
[1] 0.00079
R> sigmahat <- estim_sigma(Xsim$X, method = "LN")
[1] "k =  10"
Warning message:
In estim_sigma(Xsim$X, method = "LN") :
Since you did not specify k, k was estimated using the FactoMineR estim_ncp function
```

Both estimations of σ with MAD (0.00811) and with LN (0.0079) are close since both asymptotic frameworks are plausible: the rank is quite small in comparaison to the size of the data and the strength of the signal is important. But of course, results may differ.

Other options of `adashrink` are available:

```
adashrink <- function(X, sigma = NA, method = c("GSURE", "QUT", "SURE"),
method.optim = "BFGS", gamma.seq = seq(1, 5, by=.1), lambda0 = NA,
center = "TRUE", nbsim = 500)
```

If σ^2 is known, it can be provided. SURE (6) and GSURE (7) are optimized using the `optim` function in R with its optimization method `method.optim = "BFGS"`. Minimization over (λ, γ) is performed continously over λ and discretly over a grid for γ (defined in `gamma.seq`); the initial value for λ can be given in `lambda0` and must be in log scale (by default the median of the singular values). The argument `center = "TRUE"` is used to center the matrix X. Finally, the argument `nbsim` is used for `method = "QUT"`. Indeed, λ^{QUT} is the upper quantile of some distribution, and is evaluated empirically by a Monte Carlo simulation which generates `nbsim` times data under the null hypothesis of no signal (that is, $\mu = O$).

With `adashrink`, it is possible to enforce soft-thresholding (4) by setting `gamma.seq = 1`. Otherwise SURE and GSURE used by ATN adapt to the data. For instance they should estimate γ close to 1 when the signal is overwhelmed by noise:

```
R> Xsim <- LRsim(n = 200, p = 500, k = 100, SNR = 0.5)
R> ada.gsure <- adashrink(Xsim$X, method = "GSURE")
R> ada.gsure$gamma
[1] 1.1
```

In **denoiseR**, the function `optishrink` estimates the signal with estimators (8) and (10):

```
optishrink <- function(X, sigma = NA, center = "TRUE", method = c("ASYMPT", "LN"),
loss = c("Frobenius", "Operator","Nuclear"),  k = NA)
```

The `method = "ASYMPT"` corresponds to the asymptotic framework of (8). In this setting, the optimal shrinker is given according to different norm losses with the `loss = "Frobenius"` by default. The other arguments are the same as explained before. For instance when `method = "LN"` and no value for sigma is provided, we use (11) and infer k as well if it is not specified. The outputs are the same as those of `adashrink`.

2.4. A bootstrap approach

All the previous estimators started from the point of view of the singular values decomposition. Josse and Wager (2014) suggested an alternative approach based on the parametric bootstrap. They considered a general framework where $X \sim \mathcal{L}(\mu)$ with μ the target rank-k matrix and $\mathcal{L}$ the distribution of the noise model, and they casted the problem as follows:

$$\hat{\mu}^{(k)^*} = XB^{(k)^*} \text{ where } B^{(k)*} = \text{argmin}_B \left\{ \mathbb{E}_{X\sim\mathcal{L}(\mu)} \left[\|\mu - XB\|_2^2 \right] : \text{rank}\,(B) \leq k \right\}. \quad (12)$$

Since the signal is unknown, they suggested using a parametric bootstrap approach by taking as a proxy for μ the data X (X is simply plugged in (12) instead of μ) which results in:

$$\widehat{B}^{(k)} = \text{argmin}_B \left\{ \mathbb{E}_{\tilde{X}\sim\mathcal{L}(X)} \left[\left\| X - \tilde{X}B \right\|_2^2 \right] : \text{rank}\,(B) \leq k \right\}. \quad (13)$$

It yields an objective which is the sum of a least squared data fitting term plus a quadratic regularizer that depends on the noise model:

$$\widehat{B}^{(k)} = \text{argmin}_B \left\{ \|X - XB\|_2^2 + \left\| S^{\frac{1}{2}} B \right\|_2^2 \right\} \quad \text{with} \quad S_{jj} = \sum_{i=1}^{n} \text{Var}_{\widetilde{X}\sim\mathcal{L}(X)} \left[\widetilde{X}_{ij} \right]$$

and the solution is

$$\hat{\mu} = X(X^\top X + S)^{-1} X^\top X. \quad (14)$$

When, the noise is Gaussian as in model (1), S is equal to a diagonal matrix with elements $n\sigma^2$ and (14) can also be written as a classical singular-value shrinkage estimator:

$$\hat{\mu}_\lambda^{(k)} = \sum_{l=1}^{k} u_l \frac{d_l}{1 + \lambda/d_l^2} v_l^\top \quad \text{with} \quad \lambda = n\sigma^2.$$

However, if we have count data and consider a Binomial model as $X_{ij} \sim \frac{1}{1-\delta}\,\text{Binomial}\,(\mu_{ij},\, 1-\delta)$ with $\delta \in (0,\, 1)$, S is a diagonal matrix with row-sums of X multiplied by $\frac{\delta}{1-\delta}$. Due to the non-isotropic noise, the new estimator $\hat{\mu}_\delta$ (14) does not reduce to singular value shrinkage: its singular vectors are also modified. This characteristic is unique and implies that the estimator is better than all its competitors to recover the signal (Josse and Wager 2014). The complexity of estimator $\hat{\mu}_\delta$ is determined via δ that we estimate by cross-validation. This procedure requires to get $\hat{\mu}_\delta$ from an incomplete data set, which will be described in Section 3.

Estimator (14) requires both the rank k and the noise variance σ^2, or the rank k and δ. Consequently, as a second step, the authors suggested to remove the rank constraint and to solve for a fixed point of the proposed denoising scheme. They iterate the procedure by replacing X with the low-rank estimate $\hat{\mu} = X\widehat{B}$ obtained from the optimal B, and by solving (13) with $\hat{B} = (\hat{\mu}^\top\hat{\mu} + S)^{-1}\hat{\mu}^\top\hat{\mu}$. This estimator is called iterated stable autoencoder (ISA) and its properties are established in Josse and Wager (2014). In simulations, it presents good results to recover the signal even though slightly inferior to those of ATN in the Gaussian case. However, in comparison to all its competitors this estimator has a great advantage outside the Gaussian noise framework and in such a setting highly dominates all the others.

Josse and Wager (2014) extended their results and also used ISA to regularize correspondence analysis (CA) (Greenacre 1984, 2007). CA is a powerful method to visualize contingency tables and consists in applying an SVD on a transformation of the data, denoted M:

$$M = R^{-\frac{1}{2}} \left(X - \frac{1}{N} r c^\top \right) C^{-\frac{1}{2}}, \quad (15)$$

where $R = \text{diag}(r)$, $C = \text{diag}(c)$, N is the the total number of counts, and r and c are vectors containing the row and column sums of X. This approach will be illustrated Section 4.2.

To perform ISA we use the function `ISA` which has the following options:

```
ISA <- function (X, sigma = NA, delta = NA, noise = c("Gaussian", "Binomial"),
transformation = c("None","CA"), svd.cutoff = 0.001, maxiter = 1000,
threshold = 1e-06, nu = min(nrow(X), ncol(X)), svdmethod = c("svd", "irlba"),
center = TRUE)
```

The CA transformation (15) can be obtained by setting the argument `transformation = "CA"`; by default it is `transformation = "NONE"`. In such a case, the noise model is by default set to `noise = "Binomial"`.

If the argument `delta` is not specified, it is estimated by default by repeated learning cross-validation as detailed previously. More precisely an internal function `estim_delta` is called:

```
estim_delta (X, delta = seq(0.1, 0.9, length.out = 9),  nbsim = 10,
transformation = c("None", "CA"),  pNA = 0.10, maxiter = 1000,
threshold = 1e-08))
```

This function returns a matrix `msep` with the prediction error obtained for δ varying from 0.1 to 0.9, as well as the value of δ minimizing the MSEP in the object `delta`. The argument `pNA` indicates the percentage of missing values inserted and predicted and the argument `nbsim` the number of times this process is repeated.

ISA can also be applied with `noise = "Gaussian"` for model (1). If σ is not specified, it is estimated by default with the function `estim_sigma` using the argument `method = "MAD"`. The option `maxiter` corresponds to the maximum number of iterations of ISA whereas `threshold` is for assessing convergence (difference between two successive iterations). It is possible to specify `svdmethod = "irlba"` to use a fast SVD which is particularly useful when dealing with large matrices. In this case, the argument `nu` which corresponds to the number of singular values to compute can be specified.

`ISA` returns as the other methods the estimator of the signal in `mu.hat`, the number of non-zero singular values in `nb.eigen` and the results of the SVD of the estimator in `low.rank`. In addition, it outputs number of iterations taken by the ISA algorithm in `nb.iter`.

3. Missing values

3.1. Iterative imputation

Low rank matrix estimation using (3) or (4) has been extended for an incomplete data set (Josse and Husson 2012; Hastie, Mazumder, Lee, and Zadeh 2015) by replacing the least squares criterion by a weighted least squares

$$\text{argmin}_{\mu} \left\{ \|W \odot (X - \mu)\|_2^2 + \lambda \|\mu\|_* \right\}, \tag{16}$$

where $W_{ij} = 0$ when X_{ij} is missing and 1 otherwise and $\odot$ stands for the elementwise multiplication. In the R package **softImpute** (Hastie and Mazumder 2015), Hastie *et al.* (2015)

solved equation (16) using an iterative imputation algorithm. It provides both an estimator for the signal μ and a completed data set. So it can be seen as matrix completion or single imputation method (Schafer 1997; Little and Rubin 1987, 2002). We follow the same rationale and define iterative imputation algorithms for ATN (5) and ISA (Section 2.4). For given regularization parameters (λ, γ), *iterative ATN algorithm* performs:

1. initialization $\ell = 0$: substitute missing values with initial values such as the mean of the variables with non-missing entries, the imputed matrix is denoted X^0.
2. step $\ell \geq 1$:
   (a) perform the SVD of $X^{\ell-1}$ to estimate quantities d^ℓ and u^ℓ, v^ℓ.
   (b) compute the fitted matrix with $\hat{\mu}^\ell_{(\lambda,\gamma)_{ij}} = \sum_{l=1}^{min(n-1,p)} d_l \max\left(1 - \frac{\lambda^\gamma}{d_l^\gamma}, 0\right) u^\ell_{il} v^\ell_{jl}$ and define the new imputed data as $X^\ell = W * X + (\mathbf{1} - W) * \hat{\mu}^\ell_{(\lambda,\gamma)}$, where $\mathbf{1}$ is a matrix of size $n \times p$ filled with ones. The observed values are the same and the missing ones are replaced by the fitted values.
3. steps (2.a) and (2.b) are repeated until the change in the imputed matrix falls below a predefined threshold $\sum_{ij}(\hat{\mu}^{\ell-1}_{(\lambda,\gamma)_{ij}} - \hat{\mu}^\ell_{(\lambda,\gamma)_{ij}})^2 \leq \varepsilon$, with ε equal to 10^{-6} for example.

In the same way, we also define an *iterative ISA algorithm* where we alternate imputation of the missing values and estimation with ISA. However, the procedure is computationally involved since ISA is already an iterative algorithm. Consequently, we advocate its use only to carry-out the cross-validation technique described Section 2.4 to estimate δ.

Although, theoretical properties of iterative ATN and ISA algorithms have not yet been investigated, we expect a good recovery of the signal and of the missing values due to their ability to accurately estimate the signal in the complete case. We checked these claims using a small simulation study detailed Section 3.3.

3.2. SURE with missing values

In comparison to the R package **softImpute** (Hastie and Mazumder 2015), in addition to providing algorithms for other estimators (ATN and ISA), we suggest a solution to select the tuning parameters with missing values. More precisely, we extend the SURE (6) criterion for missing values to select (λ, γ) as follows: we first replace the first term of the sum by the number of observed values $(np - |NA|)\sigma^2$ with $|NA|$ the number of missing cells. Then, we replace the second term by the RSS on the observed values $\sum_{ij \in obs}(X_{ij} - \hat{\mu}^{\text{miss}}_{(\lambda,\gamma)_{ij}})^2$, with $\hat{\mu}^{\text{miss}}_{(\lambda,\gamma)}$, the estimator obtained from the incomplete data with the *iterative ATN algorithm* (Section 3.1). Finally, we need to define the divergence of the estimator from an incomplete data which is a difficult task. Since no explicit form is yet available, we suggest using finite differences (FD) and define SURE with missing values as:

$$\text{SURE}^{\text{miss}} = -(np - |NA|)\sigma^2 + \sum_{ij \in obs}(X_{ij} - \hat{\mu}^{\text{miss}}_{(\lambda,\gamma)_{ij}})^2 + 2\sigma^2 \text{div}^{\text{miss}}(\hat{\mu}^{\text{miss}}_{(\lambda,\gamma)}) \tag{17}$$

with

$$\text{div}^{\text{miss}}(\hat{\mu}^{\text{miss}}_{(\lambda,\gamma)}) = \sum_{ij \in obs} \frac{\hat{\mu}^{\text{miss}}_{(\lambda,\gamma)}(X_{ij} + \delta) - \hat{\mu}^{\text{miss}}_{(\lambda,\gamma)}(X_{ij})}{\delta}, \tag{18}$$

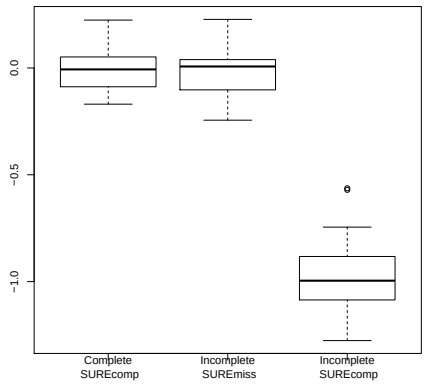

Figure 1: Boxplots of the biais for SURE: on a complete data set (left), with missing data using $\text{SURE}^{\text{miss}}$ (middle), with missing data using $\text{SURE}^{\text{comp}}$ on imputed data (right).

where δ is a small variation near machine precision. In the same way, we also suggest a GSURE criterion (7) with missing values to deal with unknown noise variance.

Note that SURE criteria with FD were also discussed in Ramani, Blu, and Unser (2008) and Deledalle, Vaiter, Fadili, and Peyré (2014). The former used it for a variety of estimators and appreciated the "black-box" aspect of the technique since no knowledge of the functional form of the estimators is required to compute it while the latter focused on the gradient of SURE FD criteria to optimize it with quasi-Newton algorithms.

To assess the validity of our proposal, we need to show that (17) is indeed an unbiased estimate of the risk of the estimator. To do so, we ran a small simulation for the soft-thresholding estimator ($\gamma = 1$) with a known variance for the noise. We consider a complete data generated according to (1) with 50 rows, 30 colums, rank $k = 5$ and SNR of 0.5. Then, for a given value of λ, we compute the *complete* estimator $\hat{\mu}_\lambda$ (4), its risk $\text{MSE}^{\text{comp}} = \text{MSE}\,(\hat{\mu}_\lambda, \mu)$ and the estimator of the risk with $\text{SURE}^{\text{comp}}$ (6). Next, we create 20% missing entries completely at random and impute the missing values with $\hat{\mu}_\lambda^{\text{miss}}$ using the *iterative ATN algorithm* described in Section 3.1. To estimate the corresponding risk, we use either $\text{SURE}^{\text{miss}}$ (17) to take imputation into account (working only on the observed values) or $\text{SURE}^{\text{comp}}$ which treats the imputed data as if they were real observations. We repeat this procedure 100 times and represent the distribution of (MSE-SURE) in Figure 1. As expected, SURE is an unbiased estimate of the risk in the complete case (left), and we see that our risk estimate (17) seems unbiased (middle), contrarily to SURE employed on the imputed data (right). Comparable results were obtained with other matrix sizes, as well as values for λ and γ.

3.3. Application to matrix completion

To perform the *iterative ATN algorithm* on an incomplete data, we use the `imputeada` function which has the following options:

```
imputeada(X, lambda = NA, gamma = NA, sigma = NA, method = c("GSURE", "SURE"),
```

```
gamma.seq = seq(1, 5, by=.1), method.optim = "BFGS", center = "TRUE",
scale = "FALSE", threshold = 1e-8, nb.init = 1, maxiter = 1000, lambda0 = NA)
```

If the argument `lambda` and `gamma` are not specified, they are estimated with `method = "GSURE"` or `method = "SURE"`. Note that contrary to the complete case, the argument `method = "QUT"` is not available and it is necessary to specify the variance of the noise when using the argument `method = "SURE"` (17). Indeed, estimates of σ with missing values are not available. The outputs are the same as for complete data, in particular the estimator of the signal obtained despite missing values is available in the object `muhat`. In addition, the function outputs a completed data set given in the object `completeObs`.

A small simulation study shows the potential of the *iterative ATN algorithm* to impute data. We used data sets that differ in terms of number of observations (n), number of variables (p) and in terms of relationships between variables. More precisely, the two first data sets are generated from model (1) with rank k respectively equal to 2 and 3 whereas the last one Parkinson data set (Stekhoven and Bühlmann 2012) is a true data set known for its nonlinear relationships between variables. Next, we inserted 20% of missing values according to the two mechanisms missing completely at random (MCAR) and missing at random (MAR) (Little and Rubin 1987, 2002) and imputed the data with the following classical methods based on:

- a Gaussian distribution assumption for the data as implemented by van Buuren and Groothuis-Oudshoorn (2011); van Buuren (2012, package **mice**). Imputation is obtained using iterative imputation with multiple regressions.
- random forests (RF) as implemented by Stekhoven and Bühlmann (2012, package **missForest**).
- the soft thresholding estimator (16) as implemented by Hastie and Mazumder (2015, package **softImpute**).
- a baseline imputation with the mean of the variables on the observed values.

Finally, we computed the errors of prediction and repeated the process 100 times. We reported in Table 1 the mean squared errors of prediction. We have selected both regularized parameters for ATN with GSURE to be as close as possible to real data analysis with no knowledge for the noise variance. Since, no method is avalaible to select the tuning parameter λ in **softImpute**, we generated a grid for λ and used its oracle value, by minimizing MSEP. The code to reproduce the simulations is available as supplementary materials.

Data	n	p	k	scale	ATN	SoftImp	RF	mice	Mean
Simulation 1	100	30	2	(1.e-5)	**1.00/1.94**	1.91/2.24	2.05/3.24	4.04/7.55	7.8/12.5
Simulation 2	41	10	3	(1.e-3)	**9/11**	19/27	68/135	89/64	261/349
Parkinson	195	22		(1)	320/1537	710/3771	**176/1503**	630/662	320/1527

Table 1: MSEP for MCAR (left)/MAR (right). Smallest MSEP row-wise is in **bold**.

The simulation results show that the proposed methods have strengths and weaknesses. The first two methods, namely ATN and SoftImpute, both perform similarly although ATN is better. This behavior was expected since ATN often improves on SoftImpute for complete data (Josse and Sardy 2015; Josse and Wager 2014). Both methods provide smallest MSEP under model (1) but struggle with non-linear relationships. Conversely, imputation with random

forests can cope with the non-linearity in the Parkinson data. However, note that despite its recent popularity (Stekhoven and Bühlmann 2012; Doove, Van Buuren, and Dusseldorp 2014; Shah, Bartlett, Carpenter, Nicholas, and Hemingway 2014), it breaks down for small sample size and MAR cases. This was already observed in Audigier, Husson, and Josse (2014) and can be explained by the inherent characteristics of the method.

4. Using denoiseR on real data

4.1. Unsupervised clustering on a tumors data

In this section, we illustrate on a tumors data set how to complement low-rank matrix estimation with unsupervised clustering. Indeed, it is common to perform k-means algorithms or hierarchical trees on the estimator given by the truncated SVD (3). The rationale is to first denoise the data to get a more stable clustering. However, if we believe in model (1), it seems more appropriate to apply the clustering on the estimators described in Section 2 rather than on (3) since the signal is better estimated with the former. Let us illustrate it on a microarray data set with 43 brain tumors of four different types defined by the standard world health organization (WHO) classification (O, oligodendrogliomas; A, astrocytomas; OA, mixed oligo-astrocytomas and GBM, glioblastomas) and 356 continuous variables corresponding to the expression data. First, we estimate the rank k, the signal with (3) and apply a hierarchical clustering on it using the using the `estim_ncp`, `PCA` and `HCPC` functions of the **FactoMineR** package. More details on these functions can be found in Lê, Josse, and Husson (2008):

```
R> data(tumors)
R> nb.ncp <- estim_ncp(tumors[ , -ncol(tumors)])
R> res.pca <- PCA(tumors, ncp = nb.ncp$ncp, quali.sup = ncol(tumors))
R> res.hcpc <- HCPC(res.pca, graph = F, consol = FALSE)
R> plot.HCPC(res.hcpc, choice = "map", draw.tree = "FALSE")
```

Then, we estimate the signal with `adashrink` and perform the same clustering algorithm:

```
R> res.ada <- adashrink(tumors[ , -ncol(tumors)], method = "SURE")
R> res.hcpc <- HCPC(as.data.frame(res.ada$mu.hat), graph = F, consol = FALSE)
R> plot.HCPC(res.hcpc, choice = "map", draw.tree = "FALSE")
```

Note that we use the `SURE` option instead of `GSURE` since the latter gave 42 non-zero singular values (the maximum number) which is not very likely.

Figures 2 shows the clustering of both approaches represented on the first two principal components. Results may conduct to a different interpretation whether using one option or the other. The truth is unknown here, but one can expect a better clustering with data denoised by ATN since simulations point to better estimation.

4.2. Regularized correspondence analysis on a Presidents data

Let's consider the `Presidents` data set, a contingency table cross-tabulating 13 US presidents (from 1940 to 2009) with 836 words used during their inaugural addresses. All the texts back

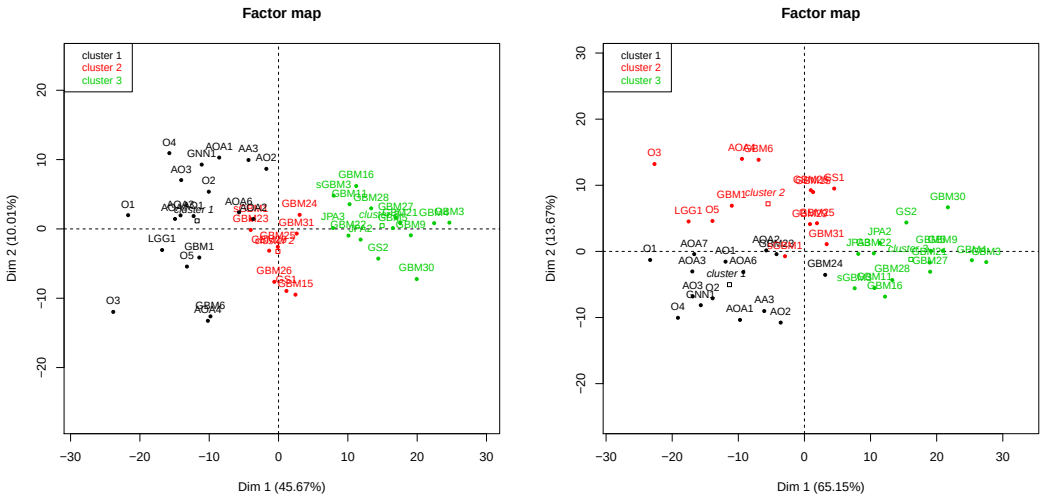

Figure 2: Clustering of the tumors data denoised by truncated SVD (left) and ATN (right).

from the speeches of George Washington in 1789 can be obtained via the public websites http://www.presidency.ucsb.edu and http://www.usa-presidents.info/union/. The texts were pre-processessed (lemmatized, keeping words that appear less than 50 times...) to get the 13×836 contingency table. The `Presidents` data set comes from the DtmVic software (Lebart 2015) specialized in the analysis of corpus data.

We perform a regularized CA using the function `ISA` and the CA transformation (15):

```
data(Presidents)
delt <- estim_delta(Presidents, transformation = "CA")
isa.ca <- ISA(Presidents, delta = delt$delta, transformation = "CA")
```

Then, the results can be visualized using the `CA` function of the **FactoMineR** package:

```
R> rownames(isa.ca$mu.hat) <- rownames(Presidents)
R> colnames(isa.ca$mu.hat) <- colnames(Presidents)
R> res.isa.ca <- CA(as.data.frame(isa.ca$mu.hat), graph = FALSE)
R> plot(res.isa.ca, title = "Regularized CA", cex = 0.7,
     selectRow = "contrib 40", invisible = "col")
R> plot(res.isa.ca, title = "Regularized CA", cex = 0.8, invisible = "row")
```

Figure 3 represents the presidents and the words on the first two CA dimensions. Briefly, rules of interpretation are the following: two presidents are close if they have the same profile of words (they over-employed or under-employed the same words); two words are close if there are associated in the same way to the presidents; a president is on the side of the words that he over-employed in comparison to the other presidents. Thus, we see that at the time of Carter, the words "soviet", "administration", "policies" were often used whereas recently "family", "parents", "child" are often mentioned. As for the tumors data in Section 2.1, it is also possible to complement the analysis by applying a clustering on the denoised data.

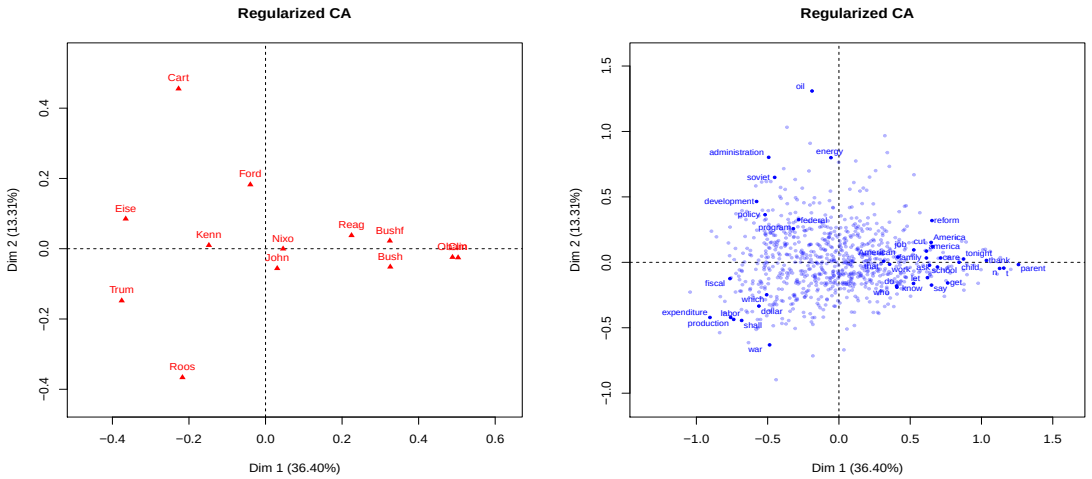

Figure 3: US Presidents from 1940 to 2009 - Words used during their inaugural speeches.

4.3. Imputation of the impact factor data

We illustrate the capability of **denoiseR** to impute on the journal impact factors data from journalmetrics.com. We use a subset of 443 journals of the same sections than Journal of Statistical Software ("Computer Science :: Software", "Decision Sciences :: Statistics, Probability and Uncertainty" and "Mathematics :: Statistics and Probability"). This data has 45 columns which correspond to three metrics recorded each year from 1999 to 2013: IPP - impact per publication (it is closed to the ISI impact factor but for three rather than two years), SNIP - source normalized impact per paper (tries to weight by the number of citations per subject field to adjust for different citation cultures) and the SJR - SCImago journal rank (tries to capture average prestige per publication). This data contains 31% of missing values. We impute it with the *ATN iterative algorithm* selecting the tuning parameters with GSURE.

```
R> data(impactfactor)
R> impactfactor[17:24,1:5]
                                SNIP_1999 IPP_1999 SJR_1999 SNIP_2000 IPP_2000
Annals of Applied Probability       1.519    0.933    1.821     1.336    0.968
Annals of Applied Statistics           NA       NA       NA        NA       NA
Annals of Probability               1.506    0.920    2.497     1.795    1.132
Annals of Software Engineering      0.529    0.176    0.203     0.690    0.299
Applied Soft Computing Journal         NA       NA       NA        NA       NA

R> ada.gsureNA <- imputeada(impactfactor, lambda = 4.46, gamma = 1.9)
R> summary(ada.gsureNA$completeObs)
```

After imputation, any analysis can be performed which is the main advantage of single imputation methods. We apply multiple factorial analysis (MFA) (Escofier and Pagès 2008; Pagès 2015) of the **FactoMineR** package. MFA is the counterpart of PCA for data with groups of variables and it allows to visualize the proximities between journals, the correlation between

metrics as well as some trajectories of the journals through the years. More details about this method can be found in Lê *et al.* (2008) and Pagès (2015). Figure 4 shows that journals on the right of the map take high values for all the metrics through the years and can be considered as "good" journals, whereas journals on the left take smaller values. In the upper part lie the journals with high scores for the SRJ measures meaning that they are prestigious. "JSS" is on the side of the good journal but when looking at its trajectory, we see that it considerably improves through the years and it is nowadays on the side of the best ones. The trajectory of the "IEEE/ACM Transactions on Networking" journal is less successful since it started correctly, was very famous in the 2002 but then slowly declined.

```
R> year=NULL; for (i in 1: 15) year= c(year, seq(i,45,15))
R> res.mfa  <-MFA(ada.gsureNA$completeObs, group=rep(3,15),  type=rep("s",15),
name.group=paste("year", 1999:2013,sep="_"),graph=F)
R>plot(res.mfa, choix = "ind", select = "contrib 15", habillage = "group", cex = 0.7)
R> points(res.mfa$ind$coord[c("Journal of Statistical Software", "Journal of the
American Statistical Association", "Annals of Statistics"), 1:2],col=2, cex=0.6)
R> text(res.mfa$ind$coord[c("Journal of Statistical Software"), 1],
res.mfa$ind$coord[c("Journal of Statistical Software"), 2],cex=1,
labels=c("Journal of Statistical Software"),pos=3, col=2)
R> plot.MFA(res.mfa,choix="var", autoLab="yes", cex=0.5,shadow=TRUE)
R> plot(res.mfa, select="Journal of Statistical Software", partial="all",
habillage="group",unselect=0.9, chrono="TRUE",xlim=c(-10,20),ylim=c(-10,20))
R> plot(res.mfa, select="IEEE/ACM Transactions on Networking", partial="all",
 habillage="group",unselect=0.9,chrono=TRUE)
```

5. Conclusion

We implemented several methods for low rank matrix estimation for complete data. The proposed development and implementation of some of these methods to handle incomplete data is novel and promising, although computationally intensive since the *iterative ATN algorithm* is performed for each cell of the data matrix to compute the divergence (18). A future version of the package will focus on improving the computational cost. Note that imputation with random forests also suffer from time problems since the method fits one random forest per variable and cycles through variables until convergence. We should also warn the user to be careful with single imputation methods. Indeed, if they apply a statistical analysis on a completed data set, the variability of the estimators will be underestimated since it does not include the variability of the prediction. Multiple imputation methods (Schafer 1997; Little and Rubin 1987, 2002) can be considered to address this issue. We finish by discussing some opportunities for further research. There appear to be many potential applications for using ISA as a method to impute count data. In addition, we would also like to define robust version of ATN with SURE with outliers cells.

6. Acknowledgement

The authors thank Achim Zeiles from providing the impact factor data and Ludovic Lebart

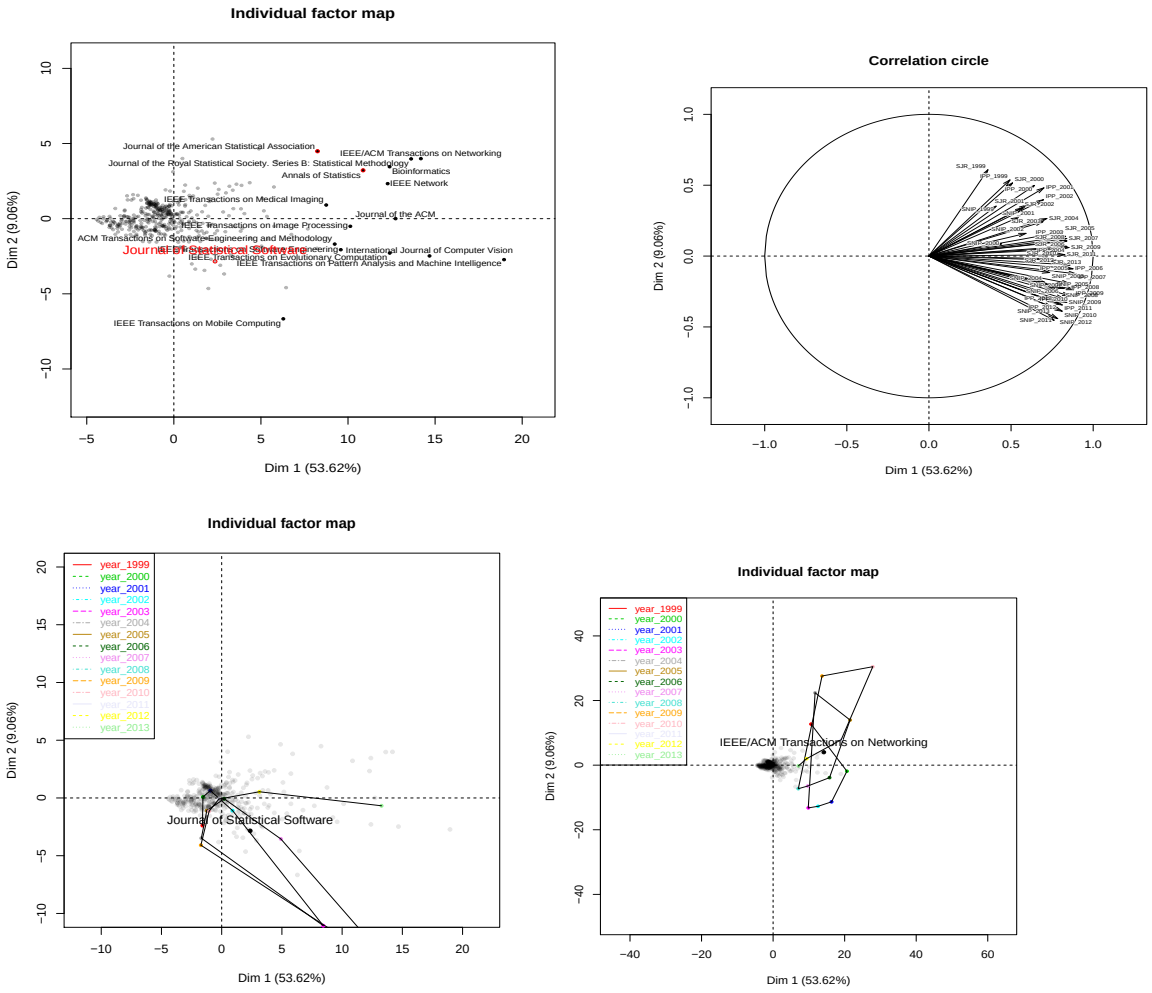

Figure 4: MFA plots. Top: journals and metrics. Bottom: trajectories of IEEE and JSS

fo the presidents data.

References

Audigier V, Husson F, Josse J (2014). "A principal components method to impute missing values for mixed data." *Advances in Data Analysis and Classification - To appear*.

Candes EJ, Sing-Long CA, Trzasko JD (2013). "Unbiased risk estimates for singular value thresholding and spectral estimators." *IEEE Transactions on Signal Processing*, **61**(19), 4643–4657.

Craven P, Wahba G (1979). "Smoothing noisy data with spline functions: estimating the correct degree of smoothing by the method of generalized cross-validation." *Numerische Mathematik*, **31**, 377–403.

Deledalle CA, Vaiter S, Fadili JM, Peyré G (2014). "Stein Unbiased GrAdient estimator of the Risk (SUGAR) for multiple parameter selection." *arXiv:1405.1164*.

Donoho DL, Johnstone IM (1994). "Ideal Spatial Adaptation via Wavelet Shrinkage." *Biometrika*, **81**, 425–455.

Doove LL, Van Buuren S, Dusseldorp E (2014). "Recursive partitioning for missing data imputation in the presence of interaction effects." *Computational Statistics & Data Analysis*, **72**, 92–104.

Eckart C, Young G (1936). "The approximation of one matrix by another of lower rank." *Psychometrika*, **1**(3), 211–218.

Efron B, Morris C (1972). "Limiting the risk of Bayes and empirical Bayes estimators. Part II: The empirical Bayes case." *Journal of the American Statistical Association*, **67**(19), 130–139.

Escofier B, Pagès J (2008). *Analyses Factorielles Simples et Multiples*. Dunod.

Gavish M, Donoho DL (2014). "Optimal Shrinkage of Singular Values." *arXiv:1405.7511v2*.

Giacobino C, Sardy S, Diaz Rodriguez J, Hengartner N (2015). "Quantile universal threshold for model selection." *arXiv:1511.05433*.

Greenacre MJ (1984). *Theory and Applications of Correspondence Analysis*. Acadamic Press.

Greenacre MJ (2007). *Correspondence Analysis in Practice, Second Edition*. Chapman & Hall.

Hastie T, Mazumder R (2015). *softImpute: Matrix Completion via Iterative Soft-Thresholded SVD*. R package version 1.4, URL http://CRAN.R-project.org/package=softImpute.

Hastie T, Mazumder R, Lee J, Zadeh R (2015). "Matrix completion and low-rank SVD via fast alternating least squares." *To appear in Journal in Machine Learning Research*.

Husson F, Josse J, Le S, Mazet J (2015). *FactoMineR: Multivariate Exploratory Data Analysis and Data Mining*. R package version 1.31.3, URL http://CRAN.R-project.org/package=FactoMineR.

Josse J, Husson F (2011). "Selecting the Number of Components in PCA Using Cross-Validation Approximations." *Computational Statististics and Data Analysis*, **56**(6), 1869–1879.

Josse J, Husson F (2012). "Handling missing values in exploratory multivariate data analysis methods." *Journal de la Société Française de Statistique*, **153 (2)**, 79–99.

Josse J, Sardy S (2015). "Adaptive shrinkage of singular values." *Statistics and Computing*, pp. 1–10.

Josse J, Wager S (2014). "Stable Autoencoding: a flexible framework for regularized low-rank matrix estimation." *arXiv:1410.8275*.

Lê S, Josse J, Husson F (2008). "FactoMineR: An R package for Multivariate Analysis." *Journal of Statistical Software*, **25**(1), 1–18.

Lebart L (2015). *Data and text mining Visualization, inference, classification*. URL http://www.dtmvic.com/.

Little RJA, Rubin DB (1987, 2002). *Statistical Analysis with Missing Data*. John Wiley & Sons series in probability and statistics, New-York.

Netflix (2009). "Netflix Challenge." URL http://www.netflixprize.com.

Pagès J (2015). *Multiple Factor Analysis with R*. Chapman & Hall/CRC.

Ramani S, Blu T, Unser M (2008). "Monte-Carlo SURE: a black-box optimization of regularization parameters for general denoising algorithms." *IEEE Trans. Image Process.*, **17**, 1540–1554.

Schafer JL (1997). *Analysis of incomplete multivariate data*. Chapman & Hall/CRC, London.

Shabalin AA, Nobel AB (2013). "Reconstruction of a low-rank matrix in the presence of Gaussian noise." *Journal of Multivariate Analysis*, **118**, 67–76.

Shah AD, Bartlett JW, Carpenter J, Nicholas O, Hemingway H (2014). "Comparison of random forest and parametric imputation models for imputing missing data using MICE: A CALIBER study." *American Journal of Epidemiology*, **179**(6), 764–774.

Stein C (1981). "Estimation of the Mean of a Multivariate Normal Distribution." *The Annals of Statistics*, **9**, 1135–1151.

Stekhoven D, Bühlmann P (2012). "MissForest - Nonparametric missing value imputation for mixed-type data." *Bioinformatics*, **28**, 113–118.

van Buuren S (2012). *Flexible Imputation of Missing Data*. Chapman & Hall/CRC, Boca Raton.

van Buuren S, Groothuis-Oudshoorn K (2011). "mice: Multivariate imputation by chained equations in R." *Journal of Statistical Software*, **45**(3), 1–67.

Verbanck M, Husson F, Josse J (2015). "Regularized PCA to denoise and visualize data." *Statistics and Computing*, **25 (2)**, 471–486.

Affiliation:

Julie Josse
Department of Statistics
Agrocampus/ INRIA Paris Orsay University
Agrocampus Ouest
E-mail: josse@agrocampus-ouest.fr
URL: http://julie.josse.com

Volume , Issue

Submitted: yyyy-mm-dd
Accepted: yyyy-mm-dd